

Parallel I/O on Highly Parallel Systems Supercomputing '95 Tutorial M6 Notes

Bill Nitzberg¹ and Samuel A. Fineberg¹

Report NAS-95-022 December 1995

Numerical Aerospace Simulation Facility
NASA Ames Research Center, M/S 258-6
Moffett Field, CA 94035-1000

Abstract

Typical scientific applications require vast amounts of processing power coupled with significant I/O capacity. Highly parallel computer systems provide floating-point processing power at low cost, but efficiently supporting a scientific workload also requires commensurate I/O performance. To achieve high I/O performance, these systems use parallelism in their I/O subsystems, supporting concurrent access to files by multiple nodes of a parallel application and striping files across multiple disks. However, obtaining maximum I/O performance can require significant programming effort. This tutorial will present a comprehensive survey of the state of the art in parallel I/O from basic concepts to recent advances in the research community. Requirements, interfaces, architectures, and performance will be illustrated using concrete examples from commercial offerings (Cray T3D, IBM SP-2, Intel Paragon, Meiko CS-2, and workstation clusters) and academic research projects (MPI-IO, Panda, PASSION, PIOUS, and Vesta).

The material covered is roughly 30% beginner, 60% intermediate, and 10% advanced.

1. MRJ, Inc., NASA Contract NAS 2-14303, Moffett Field, CA 94035-1000